CUR Algorithm for Partially Observed Matrices

Miao Xu[1], Rong Jin[2], Zhi-Hua Zhou[1*]

[1] *National Key Laboratory for Novel Software Technology*
Nanjing University, Nanjing 210093, China
Department of Computer Science and Engineering
[2] *Michigan State University, East Lansing, MI, 48824*

arXiv:1411.0860v1 [cs.LG] 4 Nov 2014

Abstract

CUR matrix decomposition computes the low rank approximation of a given matrix by using the actual rows and columns of the matrix. It has been a very useful tool for handling large matrices. One limitation with the existing algorithms for CUR matrix decomposition is that they need an access to the *full* matrix, a requirement that can be difficult to fulfill in many real world applications. In this work, we alleviate this limitation by developing a CUR decomposition algorithm for partially observed matrices. In particular, the proposed algorithm computes the low rank approximation of the target matrix based on (i) the randomly sampled rows and columns, and (ii) a subset of observed entries that are randomly sampled from the matrix. Our analysis shows the relative error bound, measured by spectral norm, for the proposed algorithm when the target matrix is of full rank. We also show that only $O(nr \ln r)$ observed entries are needed by the proposed algorithm to perfectly recover a rank r matrix of size $n \times n$, which improves the sample complexity of the existing algorithms for matrix completion. Empirical studies on both synthetic and real-world datasets verify our theoretical claims and demonstrate the effectiveness of the proposed algorithm.

Key words: Matrix approximation, CUR algorithm, matrix completion

1. Introduction

In many machine learning applications, it is convenient to represent data by matrix. Examples include user-item rating matrix in recommender system [SRJ04], gene expression matrix in bioin-

*Corresponding author. Email: zhouzh@nju.edu.cn

formatics [MD08], kernel matrix in kernel learning [WS00], document-term matrix in document retrieval [MD08], and instance-label matrix in multi-label learning [GZR+10]. An effective approach for handling big matrices is to approximate them by their low rank counterparts which can be computed and stored efficiently. Various methods have been developed for low rank matrix approximation, including truncated singular value decomposition, matrix factorization [SRJ04], matrix regression [Kol11], column subset selection [BDMI11], the Nyström method [WS00].

In this work, we will focus on the CUR algorithm for low rank matrix approximation [MD09]. It is a randomized algorithm that computes the low rank approximation for a given rectangle matrix by randomly sampled columns and rows of the matrix. Compared to other low rank approximation algorithms, CUR is advantageous in that it has (i) an easy interpretation of the approximation result because the subspace is constructed by the actual columns and rows of the target matrix [MD09], and (ii) strong (near-optimal) theoretical guarantee [BXM10, DKM06, MD08, MD09, WZ12, WZ13]. The CUR matrix decomposition algorithm has been successfully applied to many domains, including bioinformatics [MD09], collaborative filtering [MTJ11], video background modeling [MTJ11], hyperspectral medical image analysis [MMD06], text data analysis [MD08]. In the past decade, many variants of the CUR algorithm have been developed [BXM10, DKM06, MTJ11, MMD06, MD08, MD09, WZ12, WZ13].

Despite the success, one limitation with the existing CUR algorithms is that to effectively compute the low rank approximation, they require an access to the *full* matrix, a requirement that can be difficult to fulfill. For instance, in bioinformatics, it is usually too expensive to acquire the full expression information for hundreds of genes and thousands of individuals; in crowdsourcing, when both the number of workers and instances are large, it becomes impractical to request every worker to label all the instances in study ; in social network analysis, it is often the case that only part of the links between individuals can be accurately detected. In all the above cases, due to the physical or financial constraints, we only have a partial observation of the target matrix, making it difficult to apply the existing CUR algorithm.

One way to deal with the missing entries is to first compute an unbiased estimation of the target matrix based on the observed entries, and then apply the CUR algorithm to the estimated matrix. The main shortcoming of this simple method is that the unbiased estimate can be far from the target matrix when the number of observation is small, as we will show in the empirical study.

Another approach is to recover the target matrix from the observed entries by using the matrix completion technique [CCS10, CR12]. Since most matrix completion algorithms are developed only for matrices of exactly low rank, they usually work poorly for matrices of full rank [EBN11]. We note that although an adaptive sampling approach is developed in [KS13] that does apply to matrices of full rank, they use a different sampling strategy, and their bound has a poor dependence on failure probability δ (i.e. $O(1/\delta)$), which significantly limit its application when both rows and columns are randomly sampled.

In this work, we address the challenge by developing a novel CUR algorithm, named **CUR+**, for partially observed matrix. More specifically, the proposed algorithm computes a low rank approximation of matrix M based on (i) randomly sampled rows and columns from M, and (ii) randomly sampled entries from M. Unlike most matrix completion algorithms that require solving an optimization problem involving trace norm regularization [Bac08, CCS10, JY09, MHT10, TS10], the proposed algorithm only needs to solve a standard regression problem and therefore is computationally efficient. In addition, we develop a relative error bound for the proposed CUR+ that works for both low-rank and full-rank matrices. In particular, to perfectly recover a rank-r matrix of size $n \times n$, only $O(nr \ln r)$ observed entries are needed, significantly lower than $O(nr \ln^2 n)$ for standard matrix completion theory [CR12, CT10, Gro11, KMO10, Rec11] and lower than $O(nr^{3/2} \ln r)$ for adaptive algorithm for matrix recovery [KS13]. We verify our theoretical claims by empirical studies of low rank matrix approximation.

The rest of the paper is organized as follows: Section 2 briefly reviews the related work on the CUR algorithms and matrix completion; Section 3 presents the proposed algorithm and its theoretical properties. Section 4 gives our empirical study. Section 5 concludes our work with future directions.

2. Related Work

CUR matrix decomposition. CUR algorithms compute a low rank approximation of the target matrix using the actual rows and columns of the matrix [BXM10, DKM06, GZT97, GTZ97, MD08, MD09, Ste99, Tyr00, WZ12, WZ13]. More specially, let $M \in \mathbb{R}^{n \times m}$ be the given matrix and r be the target rank for approximation. A classical CUR decomposition algorithm [MD08,

MD09] randomly samples d_1 columns and d_2 rows from M, according to their leverage scores, to form matrices C and R, respectively. The approximated matrix $\widehat{M}$ is then computed as $\widehat{M} = C(C^\dagger M R^\dagger)R$, where † is the pseudoinverse. [DKM06] gives an additive error bound for the CUR decomposition, and a relative error bound, a significantly stronger result, is given in [MD08]. It stated that, with a high probability,

$$\|M - \widehat{M}\|_F \leq (1+\epsilon)\|M - M_r\|_F \tag{1}$$

where M_r is the best rank-r approximation to M, and $\|\cdot\|_F$ is the Frobenius norm of a matrix.

Various improved versions of CUR have been developed. [MTJ11] proposes a divide-and-conquer method to compute the CUR decomposition in parallel. [WZ13] proposes an adaptive CUR algorithm with much tighter error bound and much lower time complexity. In [DKM06], the authors suggest a simple uniform sampling of columns and rows for the CUR decomposition when the maximum statistical leverage scores, also referred to as incoherence measure [CR12, CT10, Rec11], is limited. In [MDMIW12], algorithms have been developed to efficiently compute the approximated values of statistical leverage scores without having to calculate the SVD decomposition of a large matrix. As we claimed in the introduction section, all the existing CUR algorithms require the knowledge of *every* entry in the target matrix and therefore cannot be applied directly to partially observed matrices. More complete list of related work on CUR can be found in [MD08, WZ13].

CUR decomposition is closely related to column subset selection problem [BDMI11, DR10, MD08], which has been studied extensively in theoretical computer science and numerical analysis communities [MD08, MD09, WZ13]. It samples multiple columns from the target matrix M and use them as the basis to approximate M, and is often viewed as special case of the CUR algorithm. A special case of column subset selection is Nyström methods, which is usually used to approximate Positive Semi-Definitive (PSD) matrix in kernel learning [WS00]. A more complete list of related Nyström methods can be found in [JYM$^+$13].

Matrix Completion. The objective of matrix completion is to fill out the missing entries of a low-rank matrix based on the observed ones. In the standard matrix completion theory, when entries are missing uniformly at random, it requires $O(nr\ln^2 n)$ observed entries to perfectly recover the target matrix under the incoherence condition [CR12, CT10, Gro11, KMO10, Rec11]. Multiple

improvements have been developed for matrix completion, either to deal with nonuniform missing entries or to develop tighter bounds under more strict coherence conditions. [KS13] developed an adaptive sensing strategy for matrix completion that removes an $\ln n$ factor from the sample complexity. In [BJ14, CBSW14], the authors study matrix completion when observed entries are not sampled uniformly at random. [NW10, RT11] generalize matrix completion to matrix regression. In [XJZ13], the authors show that the sample complexity for perfect matrix recovery can be reduced dramatically with appropriate side information.

Although it is appealing to directly combine the CUR algorithm with matrix completion to estimate a low rank approximation of a partially observed matrix, it may not work well in practice. One issue is that most matrix completion algorithms are developed for matrix of exactly low rank, significantly limiting its application to low rank matrix application. Although a few studies develop recovery bounds for matrix of full rank [EBN11, KS13], recovery errors usually deteriorate dramatically when applied to a matrix with a long tail spectrum. In addition, most matrix completion algorithms are computationally expensive, especially for large matrices, since they require, at each iteration of optimization, computing the SVD decomposition of the approximate matrix [Bac08, CCS10, JY09, MHT10, TS10]. In contrast, the proposed CUR algorithm scales to large matrix and works well for matrix of full rank.

3. CUR+ for Partially Observed Matrices

We describe the proposed CUR+ algorithm, and then present the key theoretical results for it. Due to space limitation, we postpone all the detailed analysis to the supplementary document.

3.1. CUR+ Algorithm

Let $M \in \mathbb{R}^{n\times m}$ be the matrix to be approximated, where $n \geq m$. To approximate M, we first sample uniformly at random d_1 columns and d_2 rows from M, denoted by $A = (\mathbf{a}_1, \ldots, \mathbf{a}_{d_1}) \in \mathbb{R}^{n\times d_1}$, and $B = (\mathbf{b}_1, \ldots, \mathbf{b}_{d_2}) \in \mathbb{R}^{m\times d_2}$, respectively, where each $\mathbf{a}_i \in \mathbb{R}^n$ and $\mathbf{b}_j \in \mathbb{R}^m$ is the ith row and the jth column of M respectively. We noticed that uniform sampling of rows and columns may not be the best strategy as it does not take into account the difference between individual rows and columns. Other sampling strategies, such as sampling rows/columns based

on their statistical leverage scores [MD08] and adaptive sampling [KS13, WZ12], can be more effective. We do not choose these sampling methods because they either require an access to the full matrix [MD08], introduce serious overhead in computation [WZ12], or result in significantly worse bound when matrix is of full rank [KS13]. Finally, for simplicity of discussion, we will assume $d_1 = d_2 = d$ throughout the draft even though our algorithm and analysis can easily be extended to the case when $d_1 \neq d_2$.

Let r be the target rank for approximation, with $r \leq d$. $\widehat{U} = (\widehat{\mathbf{u}}_1, \ldots, \widehat{\mathbf{u}}_r) \in \mathbb{R}^{n \times r}$ and $\widehat{V} = (\widehat{\mathbf{v}}_1, \ldots, \widehat{\mathbf{v}}_r) \in \mathbb{R}^{m \times r}$ are the first r eigenvectors of $AA^\top$ and $BB^\top$, respectively. Besides A and B, we furthermore sample, uniformly at random, entries from matrix M. Let Ω include the indices of randomly sampled entries. Our goal is to estimate a low rank approximation of matrix M using A, B, and randomly sampled entries in Ω. To this end, we need to solve the following optimization

$$\min_{Z \in \mathbb{R}^{r \times r}} \frac{1}{2} \|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)\|_F^2 \tag{2}$$

where given Ω, we define a linear operator $\mathcal{R}_\Omega(M) : \mathbb{R}^{n \times m} \mapsto \mathbb{R}^{n \times m}$ as

$$[\mathcal{R}_\Omega(M)]_{i,j} = \begin{cases} M_{i,j} & (i,j) \in \Omega \\ 0 & (i,j) \notin \Omega \end{cases}$$

Let Z_* be an optimal solution to (2). The estimated low rank approximation is given by $\widehat{M} = \widehat{U} Z_* \widehat{V}^\top$. $\widehat{M}$ can also be expressed using standard $C \times U \times R$ formulation by solving a group of linear equations. We note that (2) is a standard regression problem and therefore can be solved efficiently using the standard regression method (e.g. accelerated gradient descent [Nes03]). We refer to the proposed algorithm as **CUR+**.

3.2. Guarantee for CUR+

Before presenting the theoretical results, we first describe the notations that will be used throughout the analysis. Let $\sigma_i, i = 1, \ldots, m$ be the singular values of M ranked in descending order, and let $\mathbf{u}_i$ and $\mathbf{v}_i$ be the corresponding left and right singular vectors. Define $U = (\mathbf{u}_1, \ldots, \mathbf{u}_m)$ and $V = (\mathbf{v}_1, \ldots, \mathbf{v}_m)$. Given $r \in [m]$, partitioning the SVD decomposition of M as

$$M = U \Sigma V^\top = \begin{array}{cc} r & m-r \\ [U_1 & U_2] \end{array} \begin{bmatrix} \Sigma_1 & \\ & \Sigma_2 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix} \tag{3}$$

Table 1: Current results of sample complexity for matrix completion (including matrix regression).

Method	CUR+	[KS13]	[BJ14]	[JNS13]	[CR12, CT10, CBSW14, KMO10, Rec11]
# Observation	$nr \ln r$	$nr^{3/2} \ln r$	nr^2	$nr^{4.5} \ln n$	$nr \ln^2 n$

Let $\widetilde{\mathbf{u}}_i, i \in [n]$ be the ith row of U_1 and $\widetilde{\mathbf{v}}_i, i \in [m]$ be the ith row of V_1. The *incoherence* measure for U_1 and V_1 is defined as

$$\mu(r) = \max\left(\max_{i\in[n]} \frac{n}{r}|\widetilde{\mathbf{u}}_i|^2, \max_{i\in[m]} \frac{m}{r}|\widetilde{\mathbf{v}}_i|^2\right) \tag{4}$$

Similarly, we can have the *incoherence* measure for matrices $\widehat{U}$ and $\widehat{V}$ that include the first r eigenvectors of $AA^\top$ and $BB^\top$, respectively. Let $\widehat{\mathbf{u}}'_i, i \in [n]$ be the ith row of $\widehat{U}$ and $\widehat{\mathbf{v}}'_i, i \in [m]$ be the ith row of $\widehat{V}$. Define the incoherence measure for $\widehat{U}$ and $\widehat{V}$ as

$$\widehat{\mu}(r) \quad = \quad \max\left(\max_{i\in[n]} \frac{n}{r}|\widehat{\mathbf{u}}'_i|^2, \max_{i\in[m]} \frac{m}{r}|\widehat{\mathbf{v}}'_i|^2\right) \tag{5}$$

Define projection operators $P_U = UU^\top$, $P_V = VV^\top$, $P_{\widehat{U}} = \widehat{U}\widehat{U}^\top$, and $P_{\widehat{V}} = \widehat{V}\widehat{V}^\top$. We will use $\|\cdot\|_2$ and $\|\cdot\|_F$ respectively for the spectral norm and Frobenius norm of a matrix.

We first present the theoretical guarantee for the CUR+ algorithm when the rank of the target matrix M is no greater than r.

Theorem 1 ***(Low-Rank Matrix Approximation)*** *Assume rank$(M) \le r$, $d \ge 7\mu(r)r(t+\ln r)$, and $|\Omega| \ge 7\mu^2(r)r^2(t+2\ln r)$. Then, with a probability at least $1-5e^{-t}$, we have $M = \widehat{M}$, where $\widehat{M}$ is a low rank approximation estimated by the CUR+ algorithm.*

Remark. Theorem 1 shows that a rank-r matrix can be perfectly recovered from $2dn + |\Omega| = O(nr \ln r)$ observed entries if we set $t = \Omega(\ln r)$. In Table 1, we compare the sample complexity of the CUR+ algorithm with the sample complexity of the other matrix completion algorithms. We observe that our result significantly improves the sample complexity from previous work. We should note that unlike [KS13] where the incoherence measure is only assumed for column vectors, we assume a small incoherence measure for both row and column vectors here. It is this

stronger assumption that allows us to sample both rows and columns, leading to the improvement in the sample complexity from $O(nr^{3/2}\ln r)$ in [KS13] to $O(nr\ln r)$.

We now consider a more general case where matrix M is of full rank. Theorem 7 bounds the difference between M and $\widehat{M}$, measured in spectral norm,

Theorem 2 *Let $r \leq m$ be an integer that is no larger than m. Assume (i) $d \geq 7\mu(r)r(t+\ln r)$, and (ii) $|\Omega| \geq 7\widehat{\mu}^2(r)r^2(t+2\ln r)$. Then with a probability at least $1-3e^{-t}$*

$$\|M-\widehat{M}\|_2^2 \leq 8\sigma_{r+1}^2\left(1+2mn\right)\left(1+\frac{m+n}{d}\right).$$

As indicated by Theorem 7, when both $\mu(r)$ and $\widehat{\mu}(r)$, the incoherence measure for the first r singular/eigen vectors of M and the sampled columns/rows, are small, we have

$$\|M-\widehat{M}\|_2 \leq O\left(\sqrt{mn}\sqrt{\frac{n}{d}}\|M-M_r\|_2\right)$$

provided that $d \geq O(r\ln r)$ and $|\Omega| \geq O(r^2\ln r)$.

One limitation with Theorem 7 is that $\widehat{\mu}(r)$ is a random variable depending on the sampled columns and rows. Since $\widehat{\mu}(r)$ can be as high as n/r, $|\Omega|$, the number of observed entries required by Theorem 7, can be as large as $O(n^2)$, making it practically meaningless. Below, we develop a result that explicitly bounds $\widehat{\mu}$ with a high probability. Using the high probability bound for $\widehat{\mu}$, we are able to show that under appropriate conditions, we need at most $O(n^2/d^2)$ observed entries in order to establish a relative error bound for $\|M-\widehat{M}\|$.

To make our analysis simple, we focus on the case when M is of full rank but with skewed singular value distribution. In particular, we assume $\sigma_r \geq \sqrt{2}\sigma_{r+1}$. In order to effectively capture the skewed singular value distribution, we introduce the concept of *numerical rank* $r(M,\eta)$ [GL96] with respect to non-negative constant $\eta > 0$

$$r(M,\eta) = \sum_{i=1}^{m}\frac{\sigma_i^2}{\sigma_i^2+mn\eta}$$

Note that when $\eta = 0$, the numerical rank is equivalent to the true rank of the matrix. The larger η is , the smaller it compared to the true rank. In the following analysis, we will replace rank r with numerical rank $r(M,\eta)$.

We furthermore generalize the definition of *incoherence* measure to matrix with *numerical rank*, that is, we further define incoherence measure $\mu(\eta)$ as

$$\mu(\eta) = \max\left(\max_{1\le i\le m} \frac{m}{r(M,\eta)}|V_{i,*}\Sigma|^2, \max_{1\le i\le n} \frac{n}{r(M,\eta)}|U_{i,*}\Sigma|^2\right) \tag{6}$$

It is easy to verify that $\mu(\eta) \ge 1$. Compared to the standard incoherence measure defined in (4), the key difference is that (6) introduces singular values Σ into the definition of incoherence measure, making it appropriate for matrix of full rank.

The following two lemmas relate $r\mu(r)$ and $r\widehat{\mu}(r)$, respectively, with $r(M,\eta)\mu(\eta)$,

Lemma 1 *If we choose $\eta = \sigma_r^2/mn$, we have*

$$r\mu(r) \le 2r(M,\eta)\mu(\eta)$$

Lemma 2 *Assume that $d \ge 16(\mu(\eta)r(M,\eta)+1)(t+\ln n)$, and $\sigma_r \ge \sqrt{2}\sigma_{r+1}$. Set $\eta = \sigma_r^2/mn$. With a probability $1-4e^{-t}$, we have*

$$r\widehat{\mu}(r) \le 2r(M,\eta)\mu(\eta) + 18n\delta^2/r \quad \text{where} \quad \delta^2 = \frac{4}{d}(\mu(\eta)r(M,\eta)+1)(t+\ln n)$$

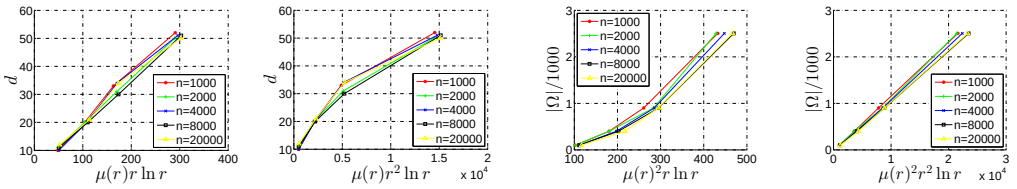
Using Theorem 7, Lemma 1 and 2, we have the result for full-rank matrix with skewed singular value distribution,

Theorem 3 *(**Full Rank Matrix Approximation**) Assume $d \ge 16(\mu(\eta)r(M,\eta)+1)(t+\ln n)$ and $\sigma_r \ge \sqrt{2}\sigma_{r+1}$. Set $\eta = \sigma_r^2/mn$. We have, with a probability $1-7e^{-t}$,*

$$\|M - \widehat{M}\|_2^2 \le 8\sigma_{r+1}^2\,(1+2mn)\left(1+\frac{m+n}{d}\right).$$

$$\text{if } |\Omega| \ge 7\left(2\mu(\eta)r(M,\eta) + 72\frac{n}{d}(\mu(\eta)r(M,\eta)+1)(t+\ln n)\right)^2(t+2\ln r) = O\left(\frac{n^2}{d^2}\right)$$

As indicated by Theorem 3, we will have a bound similar to that of Theorem 7 if $|\Omega| \ge O(n^2/d^2)$. The key difference between Theorem 7 and 3 is that in Theorem 7, the requirement for $|\Omega|$ depends on $\widehat{\mu}(r)$, a random variable depending on the sampled rows and columns. In contrast, in Theorem 3, we remove $\widehat{\mu}$ and bound $|\Omega|$ directly. We finally note that the result $|\Omega| \ge O(n^2/d^2)$ requires a large number of sampled entries for accurately estimating the low rank approximation of the target matrix. This is mostly due to the potentially loose bound for $\widehat{\mu}$. It remains an open question whether it is possible to reduce the number of observed entries for CUR-type low rank approximation.



(a) (b) (c) (d)

Figure 1: Experiment results on the synthetic data. (a)(b) plot the minimum d for perfect matrix recovery against $r \ln r$ and $r^2 \ln r$ respectively, and (c)(d) plot the minimum $|\Omega|$ for perfect matrix recovery against $r \ln r$ and $r^2 \ln r$. The results confirm the theoretical finding in Theorem 1, i.e. $d = O(r \ln r)$ and $|\Omega| = O(r^2 \ln r)$.

4. Experiments

We first verify the theoretical result in Theorem 1, i.e. the dependence of sample complexity on r and n, using synthetic data. We then evaluate the performance of the proposed CUR+ algorithm by comparing it to the state-of-the-art algorithms for low rank matrix approximation. We implement the proposed algorithm using Matlab, and all the experiments were run on a Linux server with CPU 2.53GHz and 48GB memory.

4.1. Experiment (I): Verifying the Dependence on r and n

We will verify the sample complexity result in Theorem 1, i.e. $d \geq O(r \ln r)$ and $|\Omega| \geq O(r^2 \ln r)$. We note both the requirements on d and $|\Omega|$ are independent from matrix size.

Settings. Here we study square matrices of different sizes and ranks, with n varied in $\{1,000; 2,000; 4,000; 8,000; 10,000\}$, and r varied in $\{10, 20, 30, 50\}$. For each special n and r, we search for the smallest d and $|\Omega|$ that can lead to almost perfect recovery of the target matrix (i.e. $\|M - \widehat{M}\|_F / \|M\|_F \leq 2 \times 10^{-4}$) in all 10 independent trials. To create the rank-r matrix $M \in \mathbb{R}^{n \times n}$, we first randomly generate matrix $M_L \in \mathbb{R}^{n \times r}$ and $M_R \in \mathbb{R}^{r \times n}$ with each entry of M_L and M_R drawn independently at random from $\mathcal{N}(0, 1)$, and M is given by $M = M_L \times M_R$. To create A and B, we sample uniformly at random d rows and columns. We further sample $|\Omega|$ entries from M to be partially observed. Under this construction scheme, the difference between the incoherence $\mu(r)$ for different sized matrices are relatively small (from minimum 1.4127 to maximum 2.4885). Although we will plot d and $|\Omega|$'s dependence on $\mu(r)$, we will ignore their impact in discussion of the results.

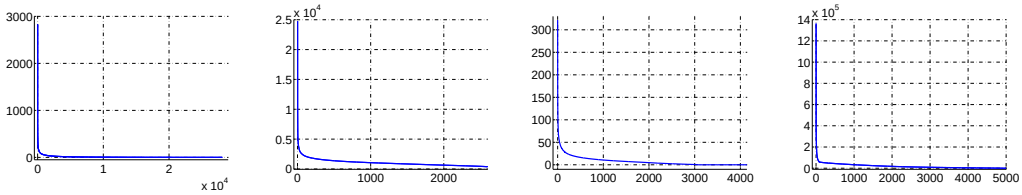


Figure 2: Singular values of real data ranked in descending order. All these four data sets are full-rank and have skewed singular value distribution.

Results. The dependence of minimal d on r and n is given in Figure 1(a) and (b), where (a) plots d against $r \ln r$ and (b) shows d versus $r^2 \ln r$. We can see clearly that d has a linear dependence on $r \ln r$. We also observed from Figure 1(a) that d is almost independent from n, the matrix size. Figure 1(c) and (d) plot the $|\Omega|$, the minimum number of observed entries, against $r \ln r$ and $r^2 \ln r$. The result in Figure 1 (d) confirms our theoretical finding, i.e. $|\Omega| \propto r^2 \ln r$.

4.2. Experiment(II): Comparison with Baseline Methods for Low Rank Approximation

We evaluate the performance of the proposed CUR+ algorithm on several benchmark data sets that have been used in the recent studies of the CUR matrix decomposition algorithm, including Enron emails $(39,861 \times 28,102)$, Dexter $(20,000 \times 2,600)$, Farm Ads $(54,877 \times 4,143)$ and Gisette $(13,500 \times 5,000)$, where each row of the matrix corresponds to a document and each column corresponds to a term/word. Detailed information of these data sets can be found in [WZ13]. All four matrices are of full rank and have skewed singular value distribution, as shown in Figure 2

Baselines. Since both the rows/columns and entries observed in the proposed algorithm are sampled uniformly at random, we only compare our approach to the standard CUR algorithm using uniformly sampled rows and columns. Although the adaptive sampling based approaches [KS13] usually yield lower errors than the standard CUR algorithm, they do not choose observed entries randomly and therefore are not included in the comparison. Let C be a set of d_1 sampled columns and R be the set of d_2 sampled rows. The low rank approximation by the CUR algorithm is given by $\widehat{M} = CZR$, where $Z \in \mathbb{R}^{d_1 \times d_2}$. Two methods are adopted to estimate Z. We first estimated Z by $Z = C^\dagger M R^\dagger$. Since this estimation requires an access to the full matrix, we refer to it as **CUR-F**. In the second method, we first construct an unbiased estimator M_e by using

the randomly observed entries in Ω, and then estimate matrix Z by $Z = C^{\dagger} M_e R^{\dagger}$. Here, the unbiased estimation M_e is given by

$$[M_e]_{i,j} = \begin{cases} \frac{mn}{|\Omega|} M_{i,j} & (i,j) \in \Omega \\ 0 & (i,j) \notin \Omega \end{cases}$$

We call this algorithm **CUR-E**. Evidently, CUR-F is expected to work better than our proposal and will provide a lower bound for the CUR algorithm for partially observed matrices.

Settings. To make our result comparable to the previous studies, we adapted the same experiment strategy as in [WZ12, WZ13]. More specially, for each data set, we set $d_1 = \alpha r$ and $d_2 = \alpha d_1$, with rank r varied in the range of $(10, 20, 50)$ and α varied from 1 to 5. To create partial observations, we randomly sample $|\Omega| = \Omega_0 = nmr^2/nnz(M)$ entries from the target matrix M, where $nnz(M)$ is the number of non-zero entries of M. We measure the performance of low rank matrix approximation by the relative spectral-norm difference $\ell_s = \|M - \widehat{M}\|/\|M - M_r\|$ which has solid theoretical guarantee according to Theorem 3. We noticed that most previous work report their results in the form of relative Frobenius norm, thus we will also show the results compared to the state-of-the-art algorithms for low rank matrix approximation measured by the relative Frobenius norm $\ell_F = \|M - \hat{M}\|_F/\|M - M_r\|_F$. Finally, we follow the experimental protocol specified in [WZ12] by repeating every experiment 10 times and reporting the mean value.

Results. Figure 3 shows the results of low rank matrix approximation for $r = 10, 20, 50$. We observe that the CUR+ works significantly better than the CUR-E method, and yields a similar performance as the CUR-F that has an access to the full target matrix M.

We observe that with larger α (i.e. increasing numbers of rows and columns), the approximation errors for CUR+ and CUR-F decrease while, to our surprise, the error of CUR-E increases significantly. This counter-intuitive result can be explained by the fact that CUR-E estimates matrix Z based on the observed entries in Ω. Since the size of Z is $d_1 \times d_2$, which increases at the rate of α^3. But on the other hand, $|\Omega|$, the number of observed entries based on which Z is estimated, remains unchanged. As a result, with increasing values of α, it becomes more and more difficult to come up with an accurate estimation of Z and consequentially a worse and worse approximation of M. We have verified this explanation in Fig 4 by simultaneously increasing the

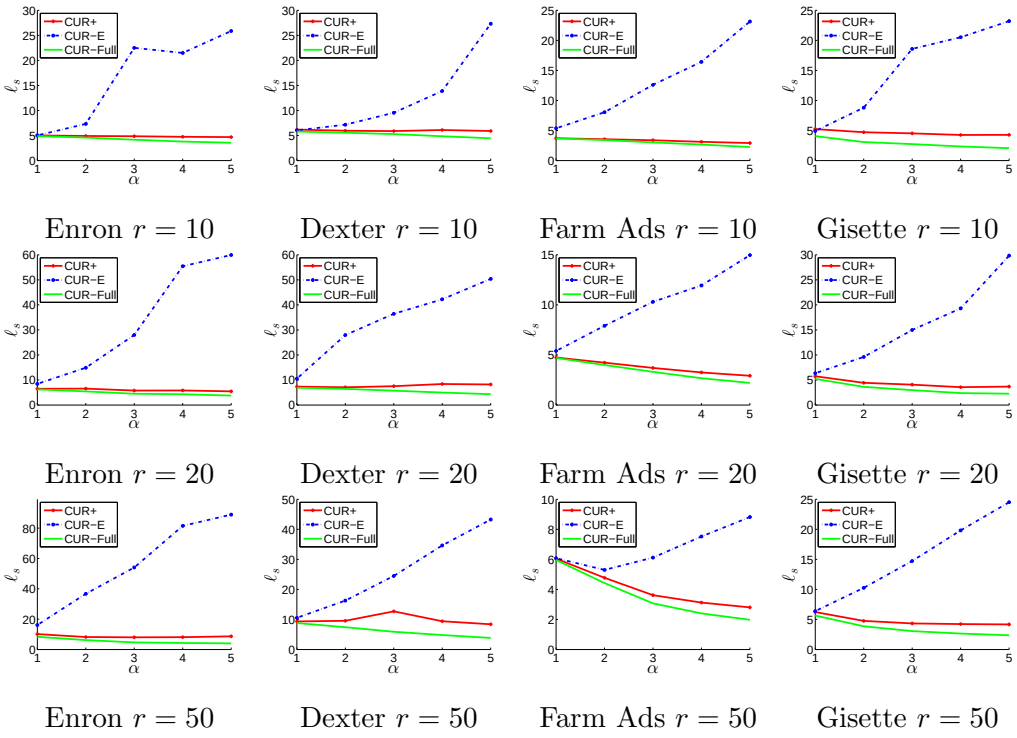


Figure 3: Comparison of CUR algorithms with the number of observed entries $|\Omega|$ fixed as $|\Omega| = \Omega_0$. The number of sampled columns and rows are set as $d_1 = \alpha r$ and $d_2 = \alpha d_1$, respectively, where $r = 10, 20, 50$ and α is varied between 1 and 5.

number of observed entries in Ω and observing that the approximation error of CUR-E decreases with increasing α, although with perturbation. It is also to our surprise that when r is increasing, the relative spectral-norm difference ℓ_s is increasing. This may due to the fact that we normalize the spectral-norm, dividing it by $\|M - M_r\|$ which decreases fast. And we observe that $\|M - \hat{M}\|$ decreases when r becomes larger and larger.

In the second experiment, we fix the number of sampled rows and columns and vary the number of observed entries from Ω_0 to $5\Omega_0$. Figure 5 shows the results for $r = 10, 20$ and 50. Again, we found that CUR+ yields similar performance as CUR-F, and performs significantly better than CUR-E, although the gap between CUR+ and CUR-E does decline with increasing number of observed entries. It is also to our surprise that for datasets Enron and Farm Ads, the approximation error of CUR+ remains almost unchanged with increasing number of observed entries. We plan to examine this unusual phenomenon in the future.

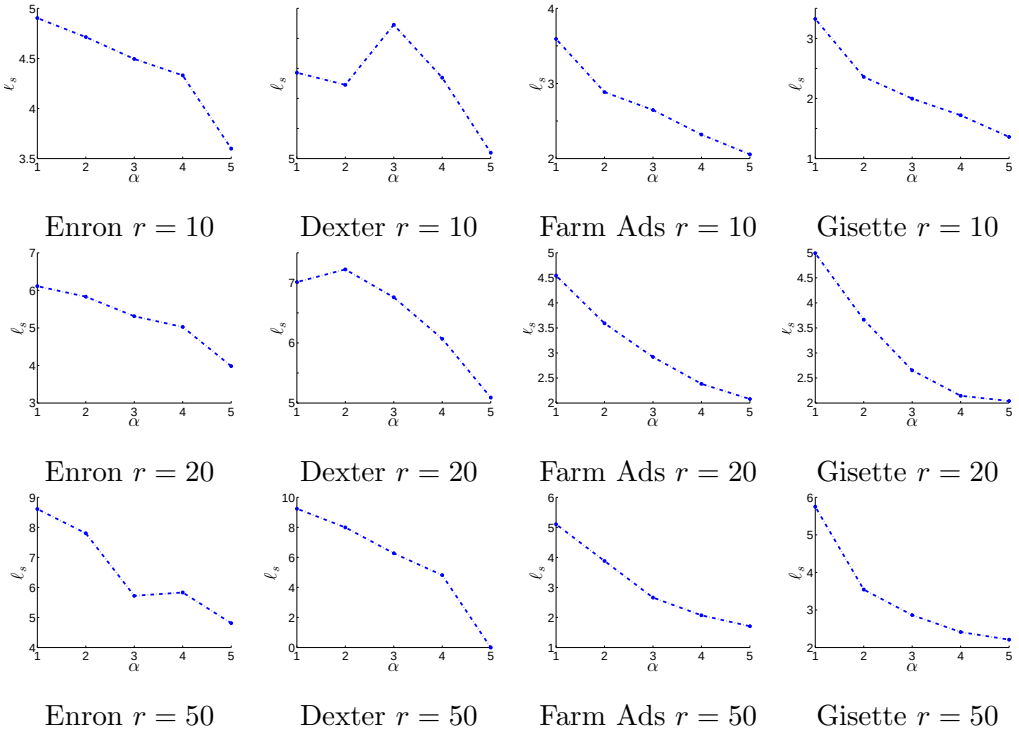


Figure 4: The results of CUR-E algorithm when $|\Omega|$ increases with α^3 for $r = 10$, 20 and 50.

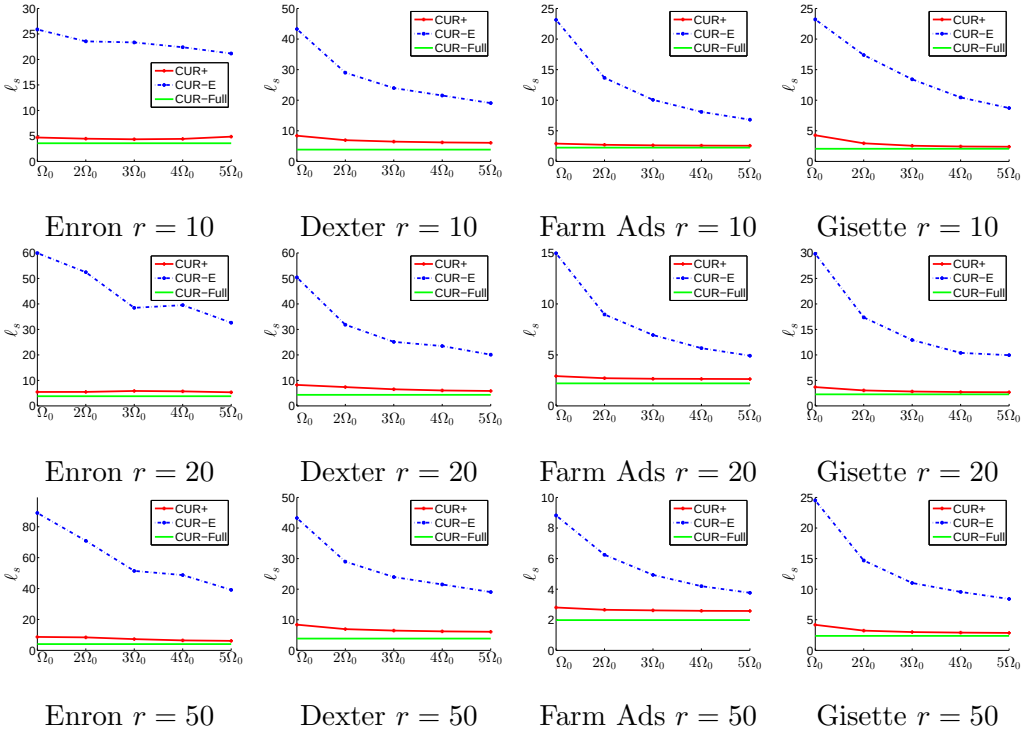


Figure 5: Comparison of CUR algorithms with the number of sampled columns (rows) fixed as $d_1 = 5r$ ($d_2 = 5d_1$), where $r = 10, 20, 50$. The number of observed entries $|\Omega|$ is varied Ω_0 to $5\Omega_0$.

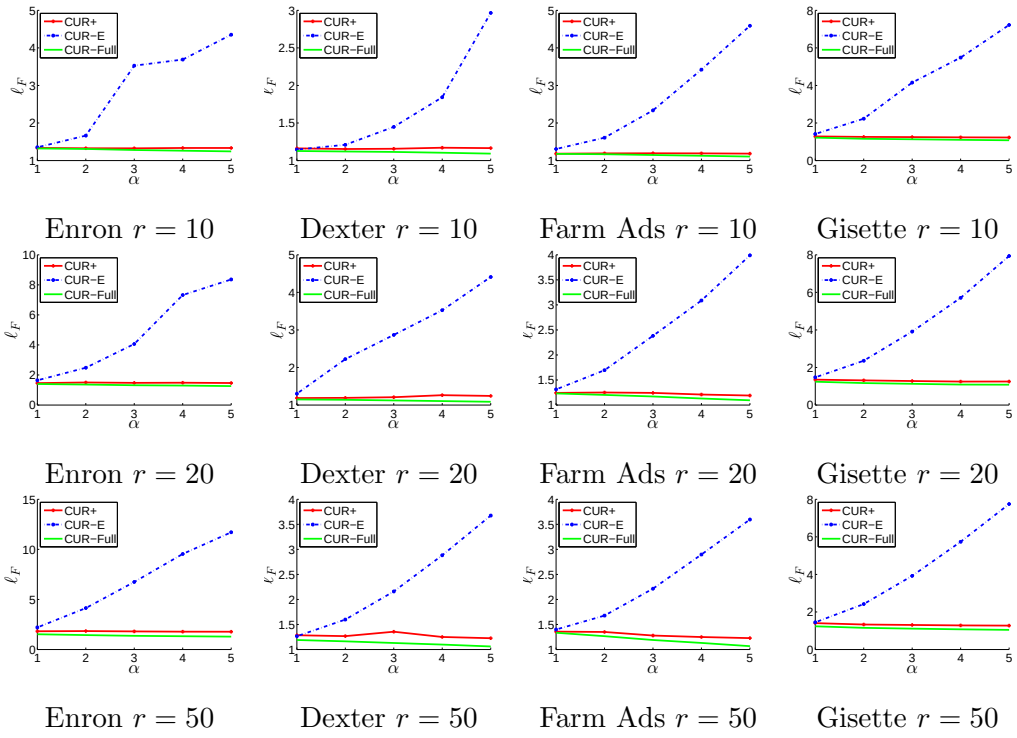


Figure 6: Comparison of CUR algorithms measured by Frobenius norm with the number of observed entries $|\Omega|$ fixed as $|\Omega| = \Omega_0$. The number of sampled columns and rows are set as $d_1 = \alpha r$ and $d_2 = \alpha d_1$, respectively, where $r = 10, 20, 50$ and α is varied between 1 and 5.

Results Measured by Frobenius Norm. Similar results on relative Frobenius norm are also reported. The results are plotted in Figure 6 when $|\Omega|$ is fixed and we vary α, and in Figure 7 when α is fixed and we vary $|\Omega|$. We can see that similar as the results measured by spectral norm, the proposed CUR+ works significantly better than the CUR-E method, and yields a similar performance as the CUR-F algorithm that has an access to the full target matrix M.

5. Conclusion

In this paper, we propose a CUR-style low rank approximation algorithm for partially observed matrix. Our analysis shows that the proposed algorithm only needs $O(nr \ln r)$ number of observed entries to perfectly recover a low-rank matrix, improving the results of the existing algorithms for matrix completion (of course under a slightly stronger condition). We also show the the spectral error bound for the proposed algorithm when the target matrix is of full rank. Empirical studies on both synthetic data and real datasets verify our theoretical claims and furthermore,

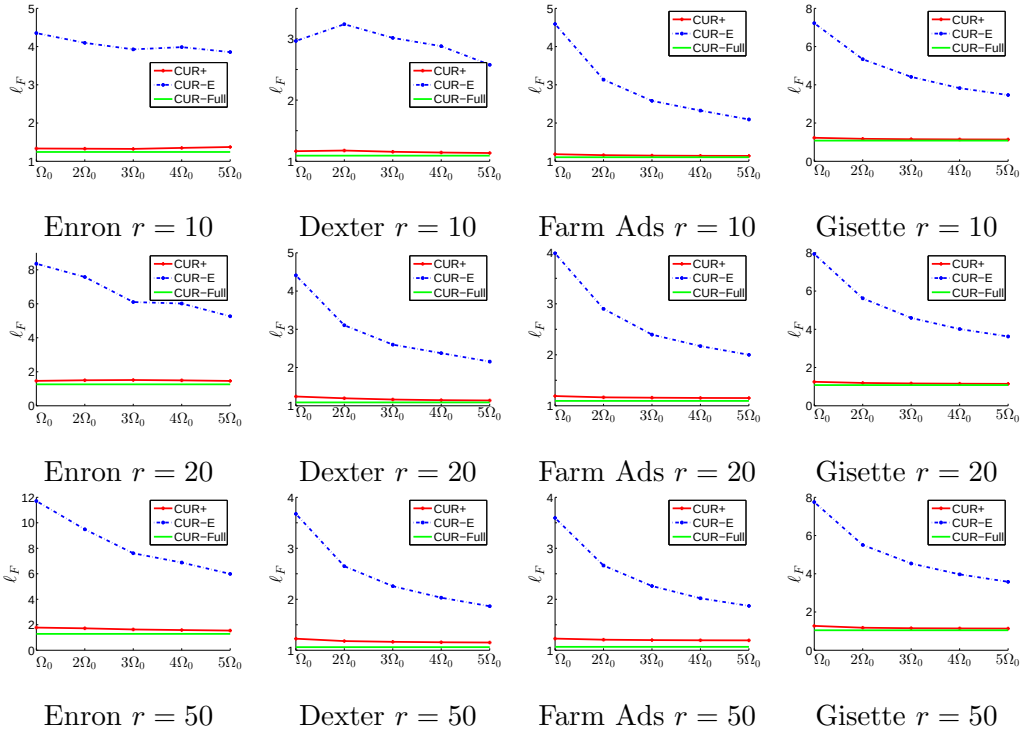


Figure 7: Comparison of CUR algorithms measured by Forbenius norm with the number of sampled columns and rows fixed as $d_1 = 5r$ and $d_2 = 5d_1$, respectively, where $r = 10, 20$ and 50. The number of observed entries $|\Omega|$ is varied from Ω_0 to $5\Omega_0$.

demonstrate that the proposed algorithm is more effective in handling partially observed matrix than the existing CUR algorithms. Since adaptive sampling has shown promising results for low rank matrix approximation [KS13], in the future, we plan to combine the proposed algorithm with adaptive sampling strategy to further reduce the error bound. We also plan to exploit the recent studies on matrix approximation/completion with non-uniform sampling and extend the CUR algorithm to the case when observed entries are non-uniform sampled.

A. Appendix

We will first give the supporting theorems we will use in the analysis. Then we will give the detailed proof of the three theorems in the paper.

A.1. Supporting Theorems

The following results are used throughout the analysis.

Theorem 4 *(Theorem 9.1 in [HMT11]) Let M be an $n \times m$ matrix with singular value decomposition $M = U\Sigma V^\top$. There is a fixed $r > 0$. Choose a test matrix $\Psi \in \mathbb{R}^{m\times d}$ and construct sample matrix $Y = M\Psi$. Partition M as in (7)*

$$M = U\Sigma V^\top = \begin{array}{cc} r & m-r \\ [U_1 & U_2] \end{array} \begin{bmatrix} \Sigma_1 & \\ & \Sigma_2 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix} \tag{7}$$

and define $\Psi_1 = V_1^\top \Psi$ and $\Psi_2 = V_2^\top \Psi$. Assuming Ψ_1 has full row rank, the approximation error satisfies

$$\|M - P_Y(M)\|_2^2 \le \|\Sigma_2\|_2^2 + \|\Sigma_2 \Psi_2 \Psi_1^\dagger\|_2^2$$

where $P_Y(M)$ projects column vectors in M in the subspace spanned by the column vectors in Y and † denotes the pseudoinverse.

Theorem 5 *(Derived From Theorem 2.2 of [Tro11]) Let $\mathcal{X}$ be a finite set of PSD matrices with dimension k (means the size of the square matrix is $k \times k$). $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ calculate the maximum and minimum eigen value respectively.*

Suppose that

$$\max_{X\in\mathcal{X}} \lambda_{\max}(X) \le B.$$

Sample $\{X_1, \ldots, X_\ell\}$ uniformly at random from $\mathcal{X}$ without replacement. Compute

$$\mu_{\max} = \ell\lambda_{\max}(\mathrm{E}[X_1]), \quad \mu_{\min} = \ell\lambda_{\min}(\mathrm{E}[X_1])$$

Then

$$\Pr\left\{\lambda_{\max}\left(\sum_{i=1}^{\ell} X_i\right) \ge (1+\rho)\mu_{\max}\right\} \le k \exp\frac{-\mu_{\max}}{B}\left[(1+\rho)\ln(1+\rho) - \rho\right] \text{ for } \rho \in [0,1)$$

$$\Pr\left\{\lambda_{\min}\left(\sum_{i=1}^{\ell} X_i\right) \le (1-\rho)\mu_{\min}\right\} \le k \exp\frac{-\mu_{\min}}{B}\left[(1-\rho)\ln(1-\rho) + \rho\right] \text{ for } \rho \ge 0$$

Theorem 6 *Let $A = S^\top HS$ and $\tilde{A} = S^\top \tilde{H} S$ be two symmetric matrices of size $n \times n$. Let $\lambda_i, i \in [n]$ and $\tilde{\lambda}_i, i \in [n]$ be the eigenvalues of A and $\tilde{A}$, respectively, ranked in descending order. Let $U_A, \tilde{U}_A \in \mathbb{R}^{n\times r}$ include the first r eigenvectors of A and $\tilde{A}$, respectively. Let $\|\cdot\|$ be any*

invariant norm. Define

$$\begin{aligned}\Delta_\lambda &= \min\left(\sqrt{2}\left(1-\frac{\lambda_{r+1}}{\lambda_r}\right), \frac{1}{\sqrt{2}}\right) \leq \frac{1}{\sqrt{2}} \\ \Delta_H &= \frac{\|H^{-1}\|\|H-\tilde{H}\|}{\sqrt{1-\|H^{-1}\|\|H-\tilde{H}\|}}\end{aligned}$$

If $\Delta_\lambda \geq \Delta_H/2$*, we have*

$$\|\sin\Theta(U_A, \tilde{U}_A)\| \leq \frac{\Delta_H}{\Delta_\lambda - \Delta_H/2}\left(1+\frac{\Delta_H\Delta_\lambda}{16}\right)$$

where

$$\Theta(X, \tilde{X}) = \arccos((X^*X)^{-1/2}X^*\tilde{X}(\tilde{X}^*\tilde{X})^{-1}\tilde{X}^*X(X^*X)^{-1/2})^{1/2}$$

defines the angle matrix between X *and* $\tilde{X}$*.*

Note that the above Theorem 6 follows directly from Theorem 4.4 and discussion in Section 5 from [Li99].

A.2. Proof of Theorem 2

We will first provide the key result for our analysis, and then bound each component of the key result, that is, first, we will show that $\|M - P_{\widehat{U}}MP_{\widehat{V}}\|_2^2$ is small; then, we will bound the strong convexity of the objective function.

The following theorem shows that the difference between M and $\widehat{M}$ is well bounded if both $\|M - P_{\widehat{U}}MP_{\widehat{V}}\|_2^2$ and the strong convexity of Eq.2 are well bounded,

Theorem 7 *Assume (i)* $\|M - P_{\widehat{U}}MP_{\widehat{V}}\|_2^2 \leq \Delta$*, and (ii) the strong convexity of the objective function is no less than* $|\Omega|\gamma$*. Then*

$$\|M - \widehat{M}\|_2^2 \leq 2\left(\Delta + \frac{\Delta}{\gamma}\right).$$

where strongly convexity is defined as,

Definition 8 *A function $f : \mathcal{D} \to \mathbb{R}$ is ξ-strongly convex w.r.t. norm $\|\cdot\|$ if f is everywhere differentiable and*

$$f(\mathbf{w}) \geq f(\mathbf{w}') + \nabla f(\mathbf{w}')(w - w') + \frac{\xi}{2}\|w - w'\|^2.$$

Then ξ is the strongly convexity of f.

Proof: Set $Z = \widehat{U}^\top M \widehat{V}$. Since $\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq \Delta$, we have

$$\|M - \widehat{U} Z \widehat{V}^\top\|_2^2 \leq \Delta,$$

implying

$$\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)\|_F^2 \leq \Delta$$

Let Z_* be the optimal solution to Eq.2. Using the strongly convexity of Eq.2, we have

$$\frac{1}{2}\gamma|\Omega|\|Z - Z_*\|_F^2 \leq \frac{1}{2}|\Omega|\Delta,$$

i.e. $\|Z - Z_*\|_F^2 \leq \Delta/(\gamma)$.

This is because $f(Z) = \frac{1}{2}\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)\|_F^2$, such that $\nabla f(Z) = \widehat{U}^T[\mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^T) - \mathcal{R}_\Omega(M)]\widehat{V}$, and $\nabla f(Z_*) = 0$

$$\begin{aligned}
\frac{|\Omega|\gamma}{2}\|Z - Z_*\|_F^2 &\leq \frac{1}{2}\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)\|_F^2 - \frac{1}{2}\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z_* \widehat{V}^\top)\|_F^2 \\
&\leq \frac{1}{2}\|\mathcal{R}_\Omega(M) - \mathcal{R}_\Omega(\widehat{U} Z \widehat{V}^\top)\|_F^2 \leq \frac{|\Omega|\Delta}{2}
\end{aligned}$$

We thus have,

$$\begin{aligned}
\|M - \widehat{M}\|_2^2 &\leq 2\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 + 2\|P_{\widehat{U}} M P_{\widehat{V}} - \widehat{U} Z_* \widehat{V}^\top\|_2^2 \\
&\leq 2\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 + 2\|P_{\widehat{U}} M P_{\widehat{V}} - \widehat{U} Z_* \widehat{V}^\top\|_F^2 \\
&\leq 2\|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 + 2\|Z - Z_*\|_F^2 \leq 2\left(\Delta + \frac{\Delta}{\gamma|}\right)
\end{aligned}$$

□

In order to bound Δ, we need the following theorem,

Theorem 9 *With a probability* $1 - 2e^{-t}$*, we have,*

$$\|M - MP_{\widehat{V}}\|_2^2 \leq \sigma_{r+1}^2 \left(1 + 2\frac{m}{d}\right)$$

and

$$\|M - P_{\widehat{U}}M\|_2 \leq \sigma_{r+1}^2 \left(1 + 2\frac{n}{d}\right)$$

provided that $d \geq 7\mu(r)r(t + \ln r)$*.*

Proof: Let $i_1, \ldots, i_d$ are the d selected columns. Define $\Psi = (\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d}) \in R^{m \times d}$, where $\mathbf{e}_i$ is the ith canonical basis. Such that we have $A = M \times \Psi$, that is, A is composed of the d selected columns of M. To utilize Theorem 4, we need to bound the minimum eigenvalue of $\Psi_1 \Psi_1^\top$, where $\Psi_1 = V_1^T \Psi \in R^{r \times d}$ is full rank. We have

$$\Psi_1 \Psi_1^\top = V_1^\top \Psi \Psi^\top V_1$$

Let $\widetilde{\mathbf{v}}_i^\top, i \in [d]$ be the ith row vector of V_1. We have,

$$\Psi_1 \Psi_1^\top = \sum_{j=1}^{d} \widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top$$

It is straightforward to show that

$$\mathrm{E}\left[\Psi_1 \Psi_1^\top\right] = \frac{d}{m} I_r$$

and

$$\mathrm{E}\left[\widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top\right] = \frac{1}{m} I_r.$$

To bound the minimum eigenvalue of $\Psi_1 \Psi_1^\top$, we need Theorem 5, where we first need to bound the maximum eigen value of $\widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top$, which is a rank-1 matrix, whose eigen value

$$\max_{1 \leq i \leq m} \lambda_{\max}(\widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top) = \max_{1 \leq i \leq m} |\widetilde{\mathbf{v}}_i|^2 \leq \mu(r)\frac{r}{m},$$

and

$$\lambda_{\max}(\mathrm{E}\left[\widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top\right]) = \lambda_{\min}(\mathrm{E}\left[\widetilde{\mathbf{v}}_{i_j} \widetilde{\mathbf{v}}_{i_j}^\top\right]) = \frac{1}{m}$$

Thus, we have,

$$\begin{aligned}\Pr\left\{\lambda_{\min}(\Psi_1\Psi_1^\top) \le (1-\delta)\frac{d}{m}\right\} &\le r\exp\frac{-d/m}{r\mu(r)/m}\left[(1-\rho)\ln(1-\rho)+\rho\right] \\ &= r\exp\frac{-d}{r\mu(r)}\left[(1-\rho)\ln(1-\rho)+\rho\right]\end{aligned}$$

By setting $\delta = 1/2$, we have,

$$\Pr\left\{\lambda_{\min}(\Psi_1\Psi_1^\top) \le \frac{d}{2m}\right\} \le r\exp\frac{-d}{7r\mu(r)} = re^{-d/[7\mu(r)r]}$$

where with $d \ge 7\mu(r)r(t+\ln r)$, we have $r\exp^{-d/[7\mu(r)r]} \le e^{-t}$, that is,

$$\Pr\left\{\lambda_{\min}(\Psi_1\Psi_1^\top) \ge \frac{d}{2m}\right\} \ge 1-e^{-t}$$

With

$$\lambda_{\min}(\Psi_1\Psi_1^\top) \ge \frac{d}{2m}$$

according to Theorem 4, we have

$$\begin{aligned}\|M - MP_{\widehat{V}}\|_2^2 &\le \|\Sigma_2\|_2^2 + \|\Sigma_2\Psi_2\Psi_1^\dagger\|_2^2 \\ &\le \sigma_{r+1}^2 + \left\|\Sigma_2\Psi_2\Psi_1^\dagger\right\|_2^2 \\ &\le \sigma_{r+1}^2 + \|\Psi_1^\dagger\|_2^2\|\Sigma_2\Psi_2\|_2^2 \\ &\le \sigma_{r+1}^2 + \frac{2m}{d}\|\Sigma_2\Psi_2\|_2^2 \\ &\le \sigma_{r+1}^2 + \frac{2m}{d}\|\Sigma_2\|_2^2\|\Psi_2\|_2^2 \\ &\le \sigma_{r+1}^2 + \frac{2m}{d}\sigma_{r+1}^2 \\ &\le \sigma_{r+1}^2\left(1+\frac{2m}{d}\right)\end{aligned}$$

- The 1st inequality is according to Theorem 4.
- The 3rd inequality is because the two facts, $\|M_1M_2\|_2 \le \|M_1\|_2 \times \|M_2\|_2$
- The 4th inequality is becuase $\|\Psi_1^\dagger\|_2 = 1/\sigma_{\min}(\Psi_1) = \sqrt{1/\lambda_{\min}(\Psi_1\Psi_1^\top)} \le \sqrt{2m/d}$
- The 6th inequality is because $\|\Sigma_2\|_2 = \sigma_{r+1}$ and $\|\Psi_2\|_2 \le \|V_2\|_2\|\Psi\|_2 = 1$

□

We then bound Δ,

Theorem 10 *With a probability* $1 - 2e^{-t}$*, we have,*

$$\Delta := \|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 \leq 4\sigma_{r+1}^2 \left(1 + \frac{m+n}{d}\right)$$

if $d \geq 7\mu(r)r(t + \ln r)$.

Proof: Using Theorem 9, we have, with a probability $1 - 2e^{-t}$

$$\begin{aligned} \|M - P_{\widehat{U}} M P_{\widehat{V}}\|_2^2 &\leq 2\|M - M P_{\widehat{V}}\|_2^2 + 2\|(M - P_{\widehat{U}} M) P_{\widehat{V}}\|_2^2 \\ &\leq 2\|M - M P_{\widehat{V}}\|_2^2 + 2\|M - P_{\widehat{U}} M\|_2^2 \\ &\leq 4\sigma_{r+1}^2 \left(1 + \frac{n+m}{d}\right) \end{aligned}$$

□

We will then bound the strong convexity of the objective function,

Theorem 11 *With a probability* $1 - e^{-t}$*, we have that* $\gamma|\Omega|$*, the strongly convexity for the objective function in (2), is bounded from below by* $|\Omega|/[2mn]$ *(that is,* $\gamma \geq 1/(2mn)$*), provided that*

$$|\Omega| \geq 7\widehat{\mu}^2(r) r^2 (t + 2\ln r)$$

Proof: To bound the strong convexity, we could instead bound the smallest eigen value of the Hessian matrix. The Hessian matrix is an $r^2 \times r^2$ matrix. Assuming the second-order derivative of the (i_1, j_1)th and (i_2, j_2)th entry of Z is the $(r(i_1 - 1) + j_1, r(i_2 - 1) + j_2)$th entry of the Hessian matrix, the Hessian matrix could be written as,

$$H = \sum_{(i,j) \in \Omega} [\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)][\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)]^T$$

To bound the minimum eigenvalue of H, we will use Lemma 5. Thus first we need to bound

$$\max_{i,j} \lambda_{\max}([\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)][\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)]^T) = \max_{i,j} |\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)|^2 \quad \leq \quad \max \|\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j\|_F^2 \leq \frac{\widehat{\mu}^2(r) r^2}{mn}$$

and

$$\begin{aligned}\lambda_{\min}\left(\mathrm{E}([\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)][\mathrm{vec}(\widetilde{\mathbf{u}}_i^\top \widetilde{\mathbf{v}}_j)]^T)\right) &= \frac{1}{mn}\lambda_{\min}\left((U\otimes V)^T \times (U\otimes V)\right) \\ &= \frac{1}{mn}\end{aligned}$$

where $\otimes$ is the Kronecker product.

Based on Theorem 5, we have

$$\Pr\left\{\lambda_{\min}(H) \leq \frac{|\Omega|}{2mn}\right\} \leq r^2 e^{\frac{-|\Omega|}{7\widehat{\mu}^2(r)r^2}}$$

Hence, with a probability $1-e^{-t}$, we have

$$\lambda_{\min}(H) \geq \frac{|\Omega|}{2mn}$$

provided that

$$|\Omega| \geq 7\widehat{\mu}^2(r)r^2(t+2\ln r)$$

□

Theorem 2 can be easily proved combining Theorems 7, 10 and 11.

A.3. Proof of Theorem 1

The following theorem allows us to replace $\widehat{\mu}(r)$ in Theorem 11 with $\mu(r)$ when the rank of M is less than or equal to r.

Theorem 12 *With a probability $1-2e^{-t}$, we have $\widehat{\mu}(r) = \mu(r)$, if $d \geq 7\mu(r)r(t+\ln r)$.*

Proof: According to Theorem 10, with a probability $1-2e^{-t}$, we have $M = P_{\widehat{U}} M P_{\widehat{V}}$, provided that $d \geq 7\mu(r)r(t+\ln r)$. Hence $P_{U_1} = P_{\widehat{U}}$ and $P_{V_1} = P_{\widehat{V}}$, which directly implies that $\mu(r) = \widehat{\mu}(r)$. □

Theorem 1 can be proved directly from Theorem 2 and Theorem 12.

A.4. Proof of Theorem 3

Define

$$H_A = \eta I + \frac{1}{mn} M M^\top, \quad \widehat{H}_A = \eta I + \frac{1}{dn} A A^\top$$

and

$$H_B = \eta I + \frac{1}{mn} M^\top M, \quad \widehat{H}_B = \eta I + \frac{1}{dm} B B^\top$$

We can have the first r eigen vector of would be H_A, because

$$\begin{aligned} H_A &= \eta I + \frac{1}{mn} M M^\top \\ &= \eta U U^T + \frac{1}{mn} U(\Sigma\Sigma^T) U^T \\ &= U(\eta I + \frac{1}{mn}\Sigma\Sigma^T) U^T \end{aligned}$$

and

$$H_A^{-1/2} = U diag(\sqrt{\frac{mn}{\sigma_1^2 + mn\eta}}, \ldots, \sqrt{\frac{mn}{\sigma_m^2 + mn\eta}}) = \sqrt{mn} U T U^T$$

where

$$T = diag(\sqrt{\frac{1}{\sigma_1^2 + mn\eta}}, \ldots, \sqrt{\frac{1}{\sigma_m^2 + mn\eta}})$$

A.4.1. Proof of Lemma 1

Proof: Just consider the maximization of the norm of rows of U, then we will have

$$\begin{aligned} \mu(\eta) &= \max_{i=1,\ldots,n} \sum_{j=1}^{m} \frac{n}{r(M,\eta)} \frac{\sigma_j^2}{\sigma_j^2 + mn\eta} U_{i,j}^2 \\ &= \max_{i=1,\ldots,n} \frac{n}{r} \sum_{j=1}^{m} r \frac{\sigma_j^2}{r(M,\eta)(\sigma_j^2 + mn\eta)} U_{i,j}^2 \\ &\geq \max_{i=1,\ldots,n} \frac{n}{r} \sum_{j=1}^{m} r \frac{a}{r} U_{i,j}^2 \\ &= a \max_{i=1,\ldots,n} \frac{n}{r} \sum_{j=1}^{m} U_{i,j}^2 \\ &= a\mu(r) \end{aligned}$$

when $\eta = \sigma_r^2/mn$, then $a \leq r/2r(M,\eta)$, then

$$\mu(r) \leq \frac{1}{a}\mu(\delta) \leq \frac{2r(M,\eta)}{r}\mu(\eta)$$

completes our proof.

□

A.4.2. Proof of Lemma 2

To this end, we need the following theorem.

Theorem 13 *With a probability* $1 - 4e^{-t}$*, we have*

$$1 - \delta \leq \lambda_k(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}) \leq 1 + \delta, \quad 1 - \delta \leq \lambda_k(H_B^{-1/2}\widehat{H}_B H_B^{-1/2}) \leq 1 + \delta, \ \forall k \in [n]$$

if

$$d \geq \frac{4}{\delta^2}(\mu(\eta)r(M,\eta) + 1)(t + \ln n)$$

Proof: It is sufficient to show the result for $\widehat{H}_A$.

Define

$$\mathcal{X} = \left\{ M_i = (H_A^{-1/2})^T \left(\frac{1}{n} M_{*,i} M_{*,i}^\top + \eta I \right) H_A^{-1/2}, i = 1, \ldots, m \right\}$$

Note that if $\mathbf{a}_i$ is the jth column of matrix M, then,

$$M_{*,i} = U\Sigma(V_{i,*})^\top$$

Thus we have

$$\begin{aligned} M_i &= mnUTU^\top(\frac{1}{n}U\Sigma V_{i,*}^\top V_{i,*}\Sigma U^\top + \eta I)UTU^\top \\ &= U\left(mT\Sigma V_{i,*}^\top V_{i,*}\Sigma T + mn\eta T^2\right)U^\top \end{aligned}$$

In this way

$$\begin{aligned}\lambda_{\max}(M_i) &\leq \lambda_{\max}(mUT\Sigma V_{i,*}^\top V_{i,*}\Sigma TU^\top) + \lambda_{\max}(mn\eta UT^2U^\top)\\ &= m|UT\Sigma V_{i,*}^\top|_2^2 + \frac{mn\eta}{\sigma_m^2 + mn\eta}\\ &\leq \mu(\eta)r(M,\eta) + 1\end{aligned}$$

(this is because $|Ax|_2^2 \leq \|A\|_2^2|x|_2^2 \leq \|A\|_F^2|x|_2^2$) and

$$\begin{aligned}\lambda_{\max}(\mathrm{E}[M_i]) &= \lambda_{\max}(U\left(T\Sigma V^\top V\Sigma T + mn\eta T^2\right)U^\top)\\ &= \lambda_{\max}(U\left(T\Sigma\Sigma T + mn\eta T^2\right)U^\top)\\ &= \frac{\sigma_1^2}{mn\eta+\sigma_1^2} + \frac{mn\eta}{mn\eta+\sigma_1^2}\\ &= 1\end{aligned}$$

So

$$\mu_{\max} = d\lambda_1(\mathrm{E}[M_i]) = d$$

we have (using Lemma 5),

$$\Pr\left\{\lambda_{\max}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \geq 1+\delta\right\} \leq n\exp\left(-\frac{d}{\mu(\eta)r(M,\eta)+1}\left[(1+\delta)\ln(1+\delta)-\delta\right]\right)$$

Using the fact that (at 0 they are the same, but the left increase faster than the right)

$$(1+\delta)\ln(1+\delta) \geq \delta + \frac{1}{4}\delta^2, \forall\delta\in[0,1],$$

we have

$$\Pr\left\{\lambda_{\max}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \geq 1+\delta\right\} \leq n\exp\left(-\frac{d\delta^2}{4(\mu r(M,\eta)+1)}\right)$$

We have the result by setting $d \geq 4(\mu(\eta)r(M,\eta)+1)(\ln n + t)/\delta^2$. Similarly, for the lower bound, we have (using Lemma 5)

$$\Pr\left\{\lambda_{\min}\left(H_A^{-1/2}\widehat{H}_A H_A^{-1/2}\right) \leq 1-\delta\right\} \leq n\exp\left(-\frac{d}{\mu(\eta)r(M,\eta)+1}\left[(1-\delta)\ln(1-\delta)+\delta\right]\right)$$

Using the fact that (by Taylor Expansion of $\ln(1-\delta)$)

$$(1-\delta)\ln(1-\delta) \geq -\delta + \frac{\delta^2}{2}$$

We have the result by setting $d \geq 2(\mu(\eta)r(M,\eta)+1)(\ln n + t)/\delta^2$. $\square$

Using Theorem 13, we will prove Lemma 2, **Proof:** To utilize Theorem 6, we rewrite H_A and $\widehat{H}_A$, as

$$H_A = H_A^{1/2} I H_A, \quad \widehat{H}_A = H_A^{1/2} D H_A^{1/2}$$

where $D = H_A^{-1/2}\widehat{H}_A H_A^{-1/2}$. According to Theorem 13, with a probability $1-2e^{-t}$, we have $\|D-I\|_2 \leq \delta$, provided that

$$d = \frac{4}{\delta^2}(\mu(\eta)r(M,\eta)+1)(t+\ln n)$$

We then compute Δ_H defined in Theorem 6 as

$$\Delta_H \leq \frac{\delta}{\sqrt{1-\delta}}$$

Because $d \geq 16(\mu(\eta)r(M,\eta)+1)(t+\ln n)$, we have

$$\frac{4}{\delta^2}(\mu(\eta)r(M,\eta)+1)(t+\ln n) \geq 16(\mu(\eta)r(M,\eta)+1)(t+\ln n)$$

that is $\delta \leq 1/2$.

Because $\sigma_r \geq \sqrt{2}\sigma_{r+1}$, we have $1/2 \leq 1-\sigma_{r+1}^2/\sigma_r^2$. Since $\delta \leq 1/2 \leq 1-\sigma_{r+1}^2/\sigma_r^2$, we have $\Delta_H \leq \sqrt{2}\delta$.

Then according to Theorem 6, we have,

$$\begin{aligned} \|\sin\Theta(U_1,\widehat{U})\|_2 &\leq \frac{\sqrt{2}\delta}{\Delta_\lambda - \sqrt{2}\delta/2}(1+\frac{\sqrt{2}\delta\Delta_\lambda}{16}) \\ &\leq \frac{\sqrt{2}\delta}{\Delta_\lambda - \sqrt{2}\delta/2}(1+\frac{1}{32}) < 3\sqrt{2}\delta \end{aligned}$$

Similarly, we have,

$$\|\sin\Theta(V_1,\widehat{V})\|_2 < 3\sqrt{2}\delta$$

Thus, with a probability $1-4e^{-t}$, we have

$$\widehat{\mu}(r) \leq \frac{2r(M,\eta)}{r}\mu(\eta) + \frac{n}{r}\|\sin\Theta(V_1,\widehat{V})\|_2^2 \leq \frac{2r(M,\eta)}{r}\mu(\eta) + \frac{18n\delta^2}{r}$$

$\square$

Theorem 3 can be proved by combining the results of Theorems 7, 11, Lemma 1 and Lemma 2.

References

[Bac08] F. Bach. Consistency of trace norm minimization. *JMLR*, 9:1019–1048, 2008.

[BDMI11] C. Boutsidis, P. Drineas, and M. Magdon-Ismail. Near optimal column-based matrix reconstruction. In *FOCS*, 2011.

[BJ14] S. Bhojanapalli and P. Jain. Universal matrix completion. In *ICML*, 2014.

[BXM10] J. Bien, Y. Xu, and M. Mahoney. Cur from a sparse optimization viewpoint. In *NIPS*, 2010.

[CBSW14] Y. Chen, S. Bhojanapalli, S. Sanghavi, and R. Ward. Coherent matrix completion. In *ICML*, 2014.

[CCS10] J.-F. Cai, E. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM J. Opti.*, 20(4):1956–1982, 2010.

[CR12] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Commun. ACM*, 2012.

[CT10] E. Candès and T. Tao. The power of convex relaxation: near-optimal matrix completion. *TIT*, 2010.

[DKM06] P. Drineas, R. Kannan, and M.W. Mahoney. Fast Monte Carlo algorithms for matrices III: Computing a compressed approximate matrix decomposition. *SIAM J. Comput.*, 36:184–206, 2006.

[DR10] A. Deshpande and L. Rademacher. Efficient volume sampling for row/column subset selection. In *FOCS*, 2010.

[EBN11] B. Eriksson, L. Balzano, and R. Nowak. High-rank matrix completion and subspace clustering with missing data. *CoRR*, 2011.

[GL96] G. Golub and C. Loan. *Matrix computations (3rd ed.)*. Johns Hopkins University Press, 1996.

[Gro11] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *TIT*, 57(3):1548–1566, 2011.

[GTZ97] S. A. Goreinov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. A theory of pseudoskeleton approximations. *Linear Algebra and Its Applications*, 261(1-3):1–21, 1997.

[GZR+10] A. Goldberg, X. Zhu, B. Recht, J.-M. Xu, and R. Nowak. Transduction with matrix completion: Three birds with one stone. In *NIPS*, 2010.

[GZT97] S. Goreinov, N. Zamarashkin, and E. Tyrtyshnikov. seudo-skeleton approximations by matrices of maximal volume. *Mathematical Notes*, 62(4):515–519, 1997.

[HMT11] N. Halko, P.-G. Martinsson, and J. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[JNS13] P. Jain, P. Netrapalli, and S. Sanghavi. Low-rank matrix completion using alternating minimization. In *STOC*, 2013.

[JY09] S. Ji and J. Ye. An accelerated gradient method for trace norm minimization. In *ICML*, 2009.

[JYM+13] R. Jin, T. Yang, M. Mahdavi, Y.-F. Li, and Z.-H. Zhou. Improved bounds for the nyström method with application to kernel classification. *TIT*, 59(10):6939–6949, 2013.

[KMO10] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *TIT*, 2010.

[Kol11] V. Koltchinskii. Low rank matrix recovery: nuclear norm penalization. In *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Springer, 2011.

[KS13] A. Krishnamurthy and A. Singh. Low-rank matrix and tensor completion via adaptive sampling. In *NIPS*, 2013.

[Li99] R.-C. Li. Relative perturbation theory: (II) eigenspace and singular subspace variations. *SIAM J. Matrix Anal. Appl.*, 20:471–492, 1999.

[MD08] M. W. Mahoney and P. Drineas. Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, 30:844–881, 2008.

[MD09] M. W. Mahoney and P. Drineas. CUR matrix decompositions for improved data analysis. *Proc. Natl. Acad. Sci. USA*, 106:697–702, 2009.

[MDMIW12] M. Mahoney, P. Drineas, M. Magdon-Ismail, and D. Woodruff. Fast approximation of matrix coherence and statistical leverage. In *ICML*, 2012.

[MHT10] R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *JMLR*, 11:2287–2322, 2010.

[MMD06] M. Mahoney, M. Maggioni, and P. Drineas. Tensor-cur decompositions for tensor-based data. In *KDD*, 2006.

[MTJ11] L. Mackey, A. Talwalkar, and M. Jordan. Divide-and-conquer matrix factorization. In *NIPS*, 2011.

[Nes03] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2003.

[NW10] S. Negahban and M. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. In *ICML*, 2010.

[Rec11] B. Recht. A simpler approach to matrix completion. *JMLR*, 12:3413–3430, 2011.

[RT11] A. Rhode and A. Tsybakov. Estimation of high dimensional low rank matrices. *Annual of Statistics*, 39(2):887–930, 2011.

[SRJ04] N. Srebro, J. Rennie, and T. Jaakkola. Maximum-margin matrix factorization. In *NIPS*, 2004.

[Ste99] G. Stewart. Four algorithms for the the efficient computation of truncated pivoted qr approximations to a sparse matrix. *Numerische Mathematik*, 1999.

[Tro11] J. Tropp. Improved analysis of the subsampled randomized hadamard transform. *Adv. Adapt. Data Anal*, 3:115–126, 2011.

[TS10] K.-C. Toh and Y. Sangwoon. An accelerated proximal gradient algorithm for nuclear norm regularized linear least squares problems. *Pacific Journal of Optimization*, 2010.

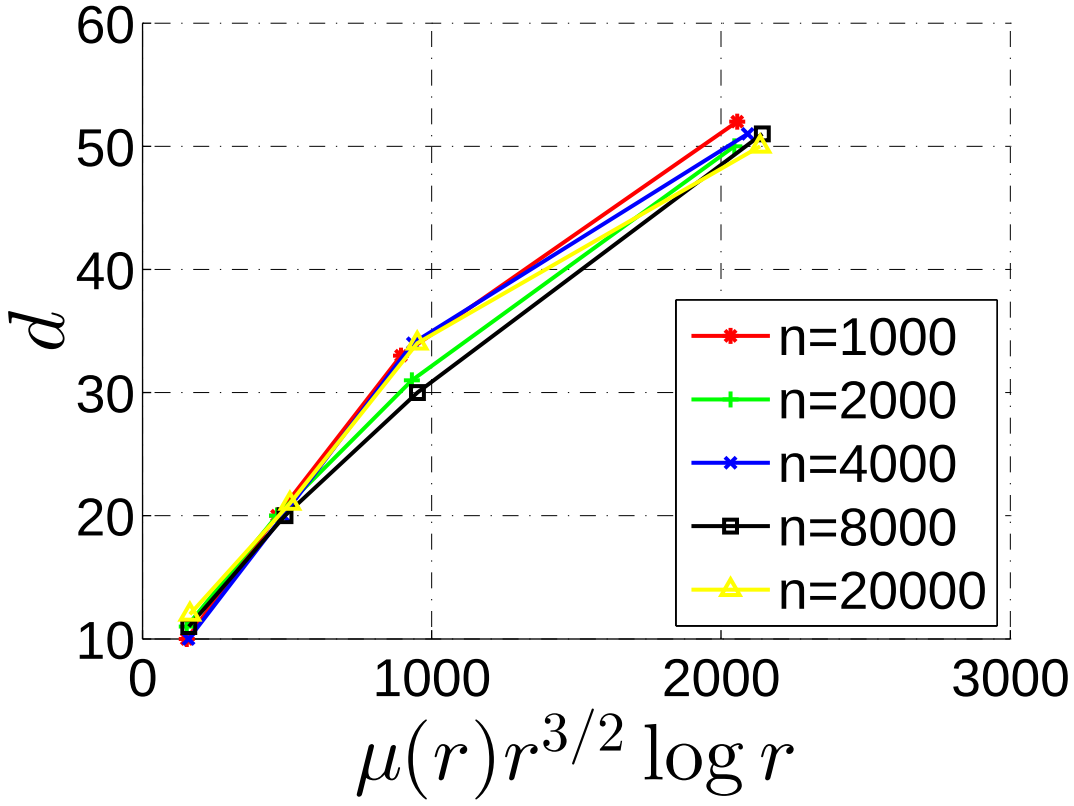
[Tyr00] E. Tyrtyshnikov. Incomplete cross approximation in the mosaic-skeleton method. *Computing*, 2000.

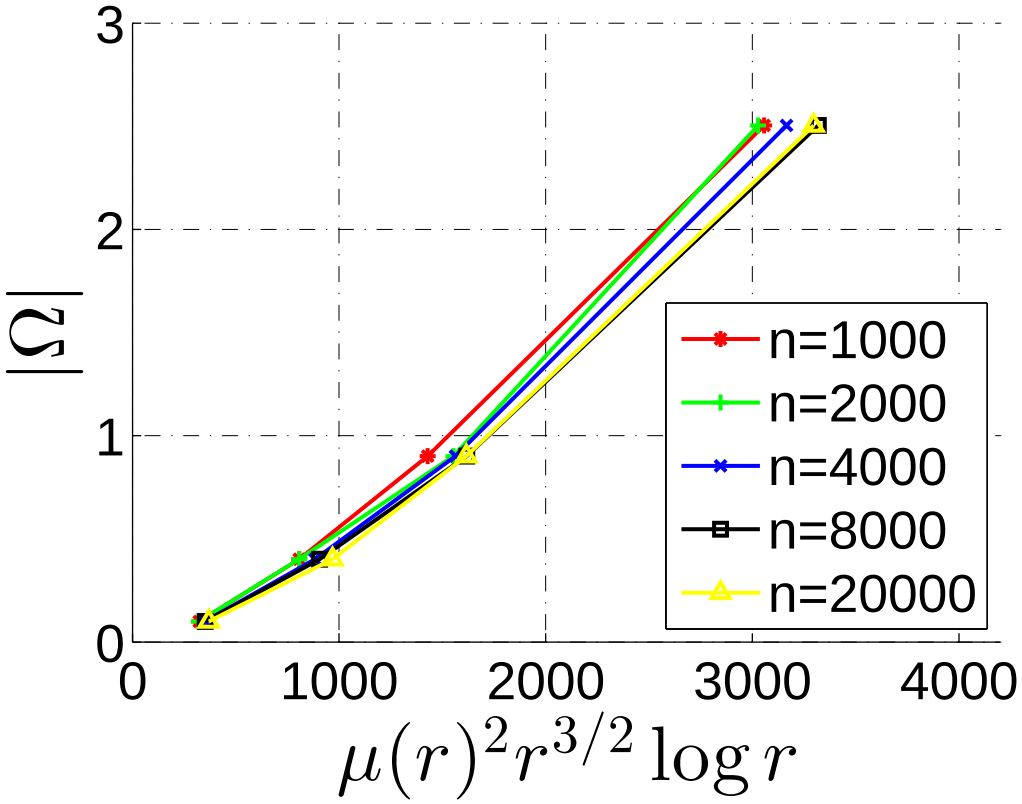
[WS00] C. Williams and M. Seeger. Using the nyström method to speed up kernel machines. In *NIPS*, 2000.

[WZ12] S. Wang and Z. Zhang. A scalable cur matrix decomposition algorithm: Lower time complexity and tighter bound. In *NIPS*, 2012.

[WZ13] S. Wang and Z. Zhang. Improving cur matrix decomposition and the nyström approximation via adaptive sampling. *JMLR*, 14(1):2729–2769, 2013.

[XJZ13] M. Xu, R. Jin, and Z.-H. Zhou. Speedup matrix completion with side information: Application to multi-label learning. In *NIPS*, 2013.


60
50
40
30
20
10
0
1000
2000
3000
d
$\mu(r)r^{3/2}\log r$
n=1000
n=2000
n=4000
n=8000
n=20000


3
2
1
0
0
1000
2000
3000
4000
$|\Omega|$
$\mu(r)^2 r^{3/2} \log r$
n=1000
n=2000
n=4000
n=8000
n=20000